

Boost Performance for Stateful Apps Deployed in Kubernetes

Lightbits Labs software-defined storage, fueled by a breadth of Intel® technologies, offers performance, resilience, and scalability for cloud-native applications in Kubernetes.



More and more businesses are turning to Kubernetes as a convenient way to deploy containerized applications. In fact, a recent report by Research and Markets with Mordor Intelligence predicts that the Kubernetes adoption rate will increase from 27 percent to 48 percent over the next five years.¹ Companies are rapidly adopting Kubernetes for good reasons: with containerized microservices, organizations can separate and scale apps and services independently for efficiency and greater resource utilization. In addition, Kubernetes provides orchestrated portability for containers, which is particularly appealing for organizations deploying cloud-native applications that support web-based and service-oriented architectures.

This portability model can break down, however, when deploying stateful applications that require persistent storage. To meet expected performance levels, apps such as MySQL, MongoDB, Redis, and Apache Spark require high input/output operations per second (IOPS) and low latency levels, together with data resiliency—particularly when running hundreds of Kubernetes clusters. Therefore, those application vendors typically recommend using direct-attached flash storage for the best performance when deploying Kubernetes on bare-metal servers. But local flash creates a problem when you're working with Kubernetes: when persistent storage is in the application server, you lose application portability, which, in turn, can result in poor flash storage utilization.

The problem for administrators, then, is how to achieve required performance levels without sacrificing Kubernetes flexibility and application portability.

One option is to use the local persistent volume feature with locally attached NVMe Express (NVMe) solid state drives (SSDs). This provides acceptable levels of performance, but it breaks the portability model of Kubernetes, and, as a result, can lead to poor storage utilization. For example, adding dedicated local flash drives to every Kubernetes server could result in 50 percent or more wasted storage.²

Alternatively, there are many Kubernetes persistent storage solutions that use Container Storage Interface (CSI), but none of them operate at the level of near-local NVMe performance without requiring a significant investment in infrastructure with proprietary hardware and drivers that lead to vendor lock-in. Proprietary drivers are often tied to specific kernel versions or require remote direct memory access (RDMA) networking, which doesn't complement the simplicity and portability philosophy of Kubernetes.

Performance with portability

With a solution from Lightbits Labs and Intel, businesses don't have to choose between performance and portability. Lightbits and Intel technologies combine to provide low-latency performance that is as close as possible to direct-attached local flash by using NVMe/TCP to support hundreds of Kubernetes clusters from

Lightbits and Intel® technology deliver a powerful platform for Kubernetes with:

- **NVM Express (NVMe)/TCP storage with local NVMe performance**
- **Kubernetes portability with persistence**
- **Rich data services**
- **Comprehensive data protection**
- **Efficient resource utilization**
- **Low total cost of ownership (TCO)**




Figure 1. Overview of Kubernetes deployed with Lightbits and Intel technologies using NVMe/TCP for persistent volumes

a single Lightbits storage cluster, all while supporting the full portability model of Kubernetes. And unlike other Kubernetes storage solutions, one Lightbits cluster can support upstream Kubernetes, Red Hat OpenShift Container Platform (OCP), VMware Tanzu, OpenStack, VMware vSphere, or bare-metal applications.

Lightbits software-defined storage, optimized for a breadth of high-performance Intel hardware technologies and low-cost SSDs, offers a complete cost-effective solution for Kubernetes running on standard TCP/IP networks and using standard NVMe/TCP drivers that are part of all major Linux distributions (see Figure 1). With the combined hardware/software solution, if a Kubernetes pod is moved to a different physical server as a result of server failure or maintenance, the same persistent volume will be attached to the new server and made available to the pod.

An Intel hardware foundation

Lightbits is optimized to provide high performance with affordability when run on commodity servers built with Intel® Xeon® Scalable processors with Intel Deep Learning Boost (Intel DL Boost), Intel® Optane™ persistent memory (PMem), Intel® Ethernet 800 Series network adapters, and low-cost quad-level cell (QLC) NAND SSDs (see Figure 2). This configuration provides consistent low-latency performance and flexibility for a wide array of configurations. It also delivers a compelling total cost of ownership (TCO) resulting from optimized flash utilization and independent scaling of storage.

The disaggregated solution uses CSI to allow containers to move between Kubernetes worker nodes on the network. But unlike other CSI-based solutions, Lightbits persistent volumes provide high-IOPS and low-latency performance that matches or exceeds direct-attached local flash,³ thanks to:

- High-performing Intel Xeon Scalable processors, available with up to 40 cores per processor


- Intel Ethernet 800 Series network adapters, providing up to 100 Gbps and using Application Device Queues (ADQ) technology, which reduces latency to the levels of RDMA, while using standard TCP/IP networking
- Intel Optane PMem, providing high-capacity memory to host metadata structures and offering large memory for high-capacity storage platforms at a lower cost

The consolidated, optimized solution can significantly outperform alternative solutions. For example, compared to Ceph, Lightbits with Intel technologies delivers 3.9x more throughput for 4K reads, 16.7x more throughput for 4K writes, 12.4x more throughput for 8K read/write workloads, 10.4x more throughput for 16K read/write workloads, and 5.6x more throughput for 32K read/write workloads.^{4,*}

The solution also includes an innovative global flash translation layer (GFTL) that enables affordable, high-capacity QLC NAND SSDs to be used for storage in place of pricier or lower-capacity alternatives with higher stated levels of endurance. The GFTL accomplishes this by relying on Intel Optane PMem to absorb writes and batch process them to the Intel QLC 3D NAND SSDs as large block sequential writes for efficiency and to help extend the lives of the drives.

In addition, the disaggregated solution lets you upgrade CPU and storage independently, as needed, providing high resource utilization because there is no need to dedicate physical flash storage to specific applications. This, combined with other operational efficiencies, helps reduce TCO compared to alternative solutions that can't offer the performance or rich feature set provided by the combination of Lightbits and Intel technologies.

^{*}See backup for workloads and configurations. Results may vary.



Intel Xeon Scalable processors

- High performance
- Storage-software optimized
- Intel Volume Management Device (Intel VMD): Enterprise-class SSD hot-plug and LED

Intel Optane technology

- Fast non-volatile write buffer and metadata
- No battery, supercaps, or services
- Large memory capacities

Intel Ethernet 800 Series with ADQ technology

- Optimized for low-latency NVMe/TCP

Figure 2. Lightbits with Intel technologies delivers performance, affordability, and high resource utilization for Kubernetes environments

Rich data services

Lightbits offers rich data services and data-protection features, including persistent volumes that can map Kubernetes availability zones to Lightbits fault domains. Lightbits also supports elastic RAID on the NVMe/TCP target servers. Replication and compression are available on a per-volume basis (one, two, or three replicas), with the ability to achieve data-reduction levels as high as 10:1 in service-provider and private-cloud environments.³ In addition, thin-volume snapshots and clones make it possible to instantly deploy high-performing development databases for efficient DevOps environments.

Get more from Kubernetes deployments

Lightbits, optimized for Intel's platform breadth of high-performance technologies and low-cost SSDs, delivers a powerful platform for Kubernetes environments. The solution offers local NVMe levels of performance with high CPU and storage utilization, consistent low latency, and rich data services with no changes to existing TCP/IP networks and no need for proprietary drivers.

The power of Intel hardware technologies paired with Lightbits software

Intel and Lightbits deliver a comprehensive solution that brings speed, scalability, and portability to the performance-hungry applications that modern businesses rely on. Intel's breadth of high-performance technologies, such as Intel Xeon Scalable processors, Intel Optane PMem, and Intel Ethernet 800 Series network adapters, is essential for delivering on the promises of NVMe/TCP to enable high-performance disaggregated storage. The integration of Lightbits with Intel technology, especially for Kubernetes environments, provides private-cloud infrastructure administrators and cloud service provider (CSP) customers with an optimized, cost-efficient, lower-TCO storage solution,² while maintaining high performance and low latency.



Learn more about Lightbits and Intel technologies

- Lightbits Labs on Kubernetes: lightbitslabs.com/kubernetes-persistent-storage/
- Intel Xeon Scalable processors: intel.com/content/www/us/en/products/details/processors/xeon/scalable.html
- Intel Ethernet 800 Series network adapters: intel.com/content/www/us/en/products/details/ethernet/800-network-adapters.html
- Intel Optane PMem: intel.com/content/www/us/en/architecture-and-technology/optane-dc-persistent-memory.html



¹ Research and Markets with Mordor Intelligence. "Application Container Market - Growth, Trends, COVID-19 Impact, and Forecasts (2021-2026)." January 2021. researchandmarkets.com/reports/4845968/application-container-market-growth-trends.

² LightOS can help lower your total cost of ownership (TCO) both for the initial purchase and over time, with greater operational efficiency. Source: Lightbits Labs. "Kubernetes and LightOS: Performance, Persistence, Simplicity." June 2021. lightbitslabs.com/ty-solutions-brief-kubernetes/.

³ Source: Lightbits Labs product page. lightbitslabs.com/product/.

⁴ Performance of Lightbits LightOS and the Intel technology platform compared to Ceph storage on Red Hat OpenShift Container Platform. Source: Testing performed by Evaluator Group and commissioned by Intel. **Configurations: Lightbits LightOS platform:** 2 x Intel Xeon Gold 6338 processors (2.00 GHz), 16 x 16 GB 3,200 megatransfers per second (MT/s) DRAM, 8 x 15.36 TB Intel SSD D5-P5316, 16 x 128 GB Intel Optane PMem 200 series, and Intel Ethernet Network Adapter E810-CQDA2 (single port used in 100 gigabit Ethernet [GbE] mode). **Ceph on Red Hat OpenShift container storage platform worker nodes:** 2 x Intel Xeon Platinum 8173M processors (2.00 GHz), 12 x 8 GB 2,666 MT/s DRAM, Intel Ethernet Network Adapter XXV710-DA2 (single port used in 25 GbE mode). **Workload description:** Software: vdbench50407. 12 worker nodes; each worker node has 12 pods running vdbench, and each pod has 1 TB persistent volume claim (PVC). All PVCs are first fully written, and then the different types of workloads (block size and read/write ratio) run for 45 minutes (nine times each), where the first 15 minutes out of the 45 are considered warm up, so the performance measurement is from minute 16 to 45.

Also see Evaluator Group white paper: Lightbits. "Run Apps up to 16X Faster: Storage Performance Comparison Lightbits vs. Ceph Storage." March 2022. lightbitslabs.com/resources/run-apps-up-to-16x-faster/

Performance varies by use, configuration and other factors. Learn more at www.intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.